

B2

**PCT**


WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification <sup>7</sup> :	A1	(11) International Publication Number:	<b>WO 00/43919</b>
G06F 17/30		(43) International Publication Date:	27 July 2000 (27.07.00)
(21) International Application Number:	PCT/US00/02190	(81) Designated States:	AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).
(22) International Filing Date:	26 January 2000 (26.01.00)	(30) Priority Data:	09/237,792 26 January 1999 (26.01.99) US
(63) Related by Continuation (CON) or Continuation-in-Part (CIP) to Earlier Application	US 09/237,792 (CIP) Filed on 26 January 1999 (26.01.99)	(71) Applicant (for all designated States except US):	APPSTREAM INC. [US/US]; 463 Addison Avenue, Palo Alto, CA 94301 (US).
(72) Inventors; and		(75) Inventors/Applicants (for US only):	RAZ, Uri [IL/US]; 36-02 Hillside Terrace, Fairlawn, NJ 07410 (US). VOLK, Yehuda [IL/IL]; 10 Rembrandt Street, Tel-Aviv (IL). MELAMED, Shmuel [IL/IL]; 35 Neve Yehushua, Ramat-Gan (IL).
(74) Agents:	KAPOUYTIAN, Ararat et al.; Morrison & Foerster LLP, 755 Page Mill Road, Palo Alto, CA 94304-1018 (US).		

(54) Title: LINK PRESENTATION AND DATA TRANSFER



(57) Abstract

A computer-implemented data access method, that may be implemented in a computer program residing on a computer-readable medium, includes receiving data at a computer from another computer prior to a user request for the data, storing the received data at the computer, and presenting a link associated with the received data using a distinguishing characteristic that differentiates the link from links to data stored only at remote computer systems. A computer system can include an interface operatively coupled to a communications link to exchange data with another computer, a display device for the presentation of data to a user, an input device to receive a user request, a memory storing executable instructions, and a processor. The processor can retrieve executable instructions from the memory to carry out aspects of the method.

BEST AVAILABLE COPY

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

## LINK PRESENTATION AND DATA TRANSFER

This is a continuation-in-part of application serial number 09/120,575 filed on July 22, 1998, entitled "Streaming Modules."

### 5 BACKGROUND INFORMATION

In a client-server environment, a client computers can communicate with a server to remotely access information stored at the server. The transfer of information between the server and client computer may be provided using standard protocols and software applications. For example, a hypertext markup language (HTML) browser application at a 10 client computer can communicate over the public Internet using TCP/IP and hypertext transfer protocols (HTTP) to receive web pages from a HTTP server. Web pages may include formatted text as well as multimedia elements, such as embedded graphics and sounds. The multimedia elements may be downloaded by the client and presented to a user by a browser application or a "plug in" browser component. Example browser applications include 15 Netscape Navigator 4.0® and Microsoft Internet Explorer 4.0™.

Browser applications used at client computers can use plug-in software to receive audio and video information using a streaming data transmission protocol. A streaming protocol allows information to be presented by a client computer as it is being received. For example, full-motion video can be sent from a server to a client as a linear stream of frames. 20 As each frame arrives at the client, it can be displayed to create a real-time full-motion video display. Audio and video streaming allows the client to present information without waiting for the entire stream to arrive at the client application. Audio and video streaming are provided by, for example, the RealAudio® and RealVideo™ applications from RealNetworks, Inc.

25 Browser applications may also make use of executable software applets to enhance the appearance of HTML-based web pages. Applets are software programs that are sent from the server to the client in response to a request from the client. In a typical applet use, HTML-based web pages include HTTP commands that cause a browser application to request an applet from a server and to begin execution of the applet. The applet may 30 thereafter interact with a user to gather and process data, may communicate data across a network, and may display results on a computer output device. Applets may be constructed

from a programming language which executes in a run-time environment provided by the browser application at the client computer. For example, the Java® programming language from Sun Microsystems, Inc., allows Java applets to be stored at a web server and attached to web pages for execution by a Java interpreter. Java Applets, may be formed from multiple 5 Java Classes. Java Classes include executable Java code that can be downloaded from a server in response to a dynamically generated request to execute the class (a module execution request). If a Java Class is not available to a Java interpreter when an executing applet attempts to access functionality provided by the Class, the Java interpreter may dynamically retrieve the Class from a server. Other programming languages, such as 10 Microsoft Visual Basic® or Microsoft Visual C++®, may also be used to create applet-like software modules, such as Microsoft ActiveX™ controls.

Downloadable applets can also be used to develop large and complex programs. For example, a complex financial program may be constructed from a collection of applets. In such a financial program, separate applets may be used to gather information from a user, 15 compute payments, compute interest, and generate printed reports. As particular program functions are required by a user, the applets associated with the required functions can be retrieved from the server. However, as the size of a software application increases, delays associated with retrieving its modules over a network likewise increase and may be unacceptable to end-users.

20 A user may repeatedly access a particular data item. For example, a web browser may be configured to display a particular "home page" when execution of the browser application begins. To reduce delays experienced when data such repeatedly accessed data is displayed, a browsers may store a copy in local "cache" memory following an initial access of the data. Caching of previously accessed data may be indicated by visual highlighting in a browser 25 display. For example, in the Netscape Version 4.0 browser, previously accessed cached data may be indicated by altering the color of a displayed HTML link. While caching by a browser, such as is performed by Netscape Version 4.0, can improve performance when retrieving previously accessed data, such a caching technique does not reduce the download time when the user initially accessed the data.

## SUMMARY

A browser or other computer application may be used to display links to data at remote computer, receive selections of links, and retrieve data from the remote computer in response to link selections. When data is received by a computer program from a remote computer, a copy can be stored in local memory or on a local hard disk drive in order to speed subsequent accesses to that data. When the data has been received in response to a user's request for the data, such as by a user clicking on a displayed link associated with the data, a browser history function may alter the color or other characteristic of the displayed link to indicate that the associated data has been previously retrieved and may be locally stored. Data may also be received from a remote computer and locally stored prior to a user request for the data. For example, data may be received by caching software that autonomously requests and receives data associated with links in a HTML document so that when a user request for the data is received, it can be quickly received from the cache. A user may be unaware that such autonomously retrieved data is locally stored. Advantages may be gained by increasing user awareness of locally stored data. For example, by identifying data that has been autonomously retrieved from a remote server and locally stored, a user can be made aware of data that can be quickly presented, and consequently, the user may be more likely to view such data.

In general, in one aspect, the invention features a computer-implemented data access method. The method includes receiving data at a computer from another computer prior to a user request for the data, storing the received data at the computer, and presenting a link associated with the received data using a distinguishing characteristic that differentiates the link from links to data stored only at remote computer systems.

In general, in another aspect, the invention features a computer program residing on a computer-readable medium. The program includes instructions for causing a computer to receive data from another computer prior to a user request for the data, to store the received data, and, prior to a user request for the data, to present a link associated with the received data using a distinguishing characteristic that differentiates the link from links to data stored only at remote computer systems.

In general, in another aspect, the invention features a computer system including an interface operatively coupled to a communications link to exchange data with another computer, a display device for the presentation of data to a user, an input device to receive a

user request, a memory storing executable instructions, and a processor. The processor is operatively coupled to the interface, the image display, the input device, and the memory. The executable instructions configure the processor to receive data over the interface from another computer prior to a receipt of a request at the input device for the data, to store the received data, and, prior to a user request for the data, to present a link associated with the received data on the display device using a distinguishing characteristic that differentiates the link from links to data stored only at remote computer systems.

Implementations may include one or more of the following features. Other data may be receiving in response to a user request, stored, and presented using a different distinguishing characteristic. A link to data not stored at the computer may be presented using another different distinguishing characteristic. Presenting using distinguishing characteristics may include presenting using different fonts, colors, graphical user interface pointer shapes, and other differences. A link may be a pointer to data at another computer. A user request may include a selection of the link, and the stored received data may be presented in response to the user request. The received data may include a program module that is presented through execution of the program module, or text that is presented by displaying the text on a video display. The other computer may include a data storage storing a collection of executable modules associated with an application. The received data may include an executable module selected from the collection, and the computer may provide an execution environment for executing the application modules. The method may also include forming a module set that includes the data. The module set may be formed by selecting a sequence of modules from the collection in accordance with predetermined criteria. The sequence of modules may then be streamed to the computer from the other computer. The module sequence may be selected in accordance with a streaming control database at the other computer. The link presentation may be altered after the presentation of data associated with the link.

In general, in another aspect, the invention features a computer system having a processor operatively interconnected to a memory, a network communications device, a graphical display device, a data storage device, and a user input device. The computer system includes a graphical user interface having interface objects that are user-selectable to present data associated with each object to a user. Each interface object may have a visual characteristics identifying the location of the objects associated data. The visual

characteristics that can be associated with an interface object include a characteristic indicating that an object's associated data is stored locally at the computer system and has not been previously presented to the user and a different characteristic indicating that an object's associated data is stored locally at the computer system and has been previously 5 presented to the user. The data stored locally at the computer system may include data received from another computer system preceding a user's selection of an interface object associated with the data from the other computer system.

The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Implementations may provide advantages 10 such as facilitating access to localized data without requiring user location input. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

#### DESCRIPTION OF DRAWINGS

- Fig. 1 illustrates a computer network.  
15 Fig. 2 illustrates computer software application modules.  
Fig. 3 is a directed graph, according to the invention.  
Fig. 4 illustrates a server and a client, according to the invention.  
Figs. 5A – 5E illustrate application code components, according to the invention.  
Fig. 6 depicts a browser interface, according to the invention.

20 **DETAILED DESCRIPTION**

A computer system, such as an industry-standard Intel® x86 compatible personal computer, an Apple Macintosh® computer, or other computer system includes both software and hardware resources. Hardware resources may include a processor, memory, a hard disk drive, an output display device, input devices such as a keyboard and a mouse, a network 25 communications device such as a modem or local area network interface. Software resources may include an operating system providing a graphical user interface and application programs. Such systems can function as client and server computers in a networked configuration.

Referring to Fig. 1, a wide area network 100 is shown. In the network 100, a client 30 computer 101 can communicate with a server computer 102 by sending data over links 103 and 104 to a data network 130. The data network 130 may include multiple nodes 131-134

that can route data between the client 101 and the server 102. The client computer 101 may transmit and receive data using the TCP/IP, HTTP, and other protocols. For example, the client 101 may use the HTTP protocol to request web pages from the server 102.

Web pages and multimedia data sent from the server 102 to the client 101 may have a natural linear sequence associated with them. The natural sequence of video data may be the linear order of video frames while the natural sequence of text may be the order in which pages of text are arranged in a document. Data having a natural linear sequence can be streamed from a server to a client to minimize download delays. In a streaming system, while earlier items in a liner sequence are being processed and/or displayed, subsequent items may be downloaded to the client computer. When processing and/or display of an item is complete, processing or display of a fully received "streamed" item may quickly begin. Since receipt of a streamed item is fully or partially complete when the item is requested, a user or client application requesting the streamed item will perceive a reduced downloading delay. For example, if the first page of a document is retrieved by a user, the second page can be downloaded while the first page is being read. If the user continues reading at the second page of the document, that page will then be available at the client, such as in a cache area on a hard disk drive, and can be read without additional downloading delay.

Software execution may not follow a predictable natural linear order. Software may include jump statements, break statements, procedure calls, and other programming constructs that cause abrupt transfers of execution among sections of executing code. The execution path that is traversed during the processing of interrelated code modules (such as code segments, code classes, applets, procedures, and code libraries), will often be non-linear, user dependent, may change with each execution of the application program, and may change depending on the state of various data items. Although a natural order may be lacking, an advantageous order may be determined in which to stream modules. The order may be determined using criteria that is independent of the computer's internal architecture or internal operating system (execution environment) considerations.

Referring to Fig. 2, a software application 200 may include multiple modules "A" through "H." Modules "A" through "H" may be Java Classes, C++ procedure libraries, or other code modules that can be stored at a server. Some of the modules "A" through "H" may also be stored at the client computer, such as in a hard disk drive cache or as part of a software library stored at the client computer. When a client computer begins execution of

the application 200, a first module, such as module "A," may be downloaded from the server and its execution at the client 410 may begin. As module "A" is being processed, the programming statements contained therein may branch to, for example, module "E." If Module "E" is not already resident at the client, the execution of module "A" can be suspended, module "E" can be retrieved from the server, and then the execution of module "E" code may begin. In such a scenario, a user will experience a module download delay associated with retrieving module "E" from the server.

To minimize module download delays experienced by a user, module "E" may be transparently streamed from a server to the client computer. Transparent streaming allows future module use to be predicted and modules to be downloaded while other interrelated modules "A" are executing. Referring to Fig. 4, an exemplary software architecture 400 providing transparent streaming is shown. The software architecture 400 includes a server 401 having a database 403 of stored software modules. The server 401 can transparently transmit a stream of software modules 405 over a communications link to a client computer 410. The communication link may be an analog modem connection, a digital subscriber line connection, a local area network connection, or any other type of data connection between the server 401 and client 410. As particular software modules are being executed at the client 410, additional modules are sent from the server 401 to the client 410. In a dynamic streaming implementation, the order in which modules are streamed between the server and client may be altered based on the particular client computer 410 being served, based on the user of the client computer, and based on other dynamically determined factors.

Referring to Fig. 3, the execution order of application modules "A" through "H" may resemble a directed graph 300 rather than a linear sequence of modules. For example, as illustrated by the graph 300, after module "A" is executed, execution can continue at module "B," "D," or "E." After module "B" is executed, execution can continue at module "C" or "G." The execution path may subsequently flow to additional modules and may return to earlier executed modules.

The server 401 can use streaming control information 402 to determine the order in which to stream modules from the server 401 to the client 410. The streaming control information 402 can include, for example, a predicted execution flow between software modules such as that represented by the directed graph 300. As downloaded modules are executed by the client 410, the client may send control data 415 to the server 401 to

dynamically update and alter the order in which modules are streamed from the server 401 to the client 410. Control data 415 may be used to request particular modules from the server 401, to send data regarding the current execution state of the application program, to detail the current inventory of modules residing in the client's local storage 411, and to report user input selections, program execution statistics, and other data derived regarding the client computer 410 and its executing software.

The sequence of modules sent in the stream 405 from the server 401 to the client 410 can be determined using a streaming control file 402. The streaming control file 402 includes data used by the server to predict modules that will be needed at the client 410. In a graph-based implementation, the control file 402 may represent modules as nodes of a directed graph. The control file 402 may also represent possible execution transitions between the modules as vertices ("edges") interconnecting the nodes. Referring to Table 1, in a weighted graph implementation, the streaming control file 402 may include a list of vertices representing all possible transitions between modules. For example, Table 1 lists vertices representing all possible transitions between the modules "A" through "H" of graph 300 (Fig. 3). Each vertex in Table 1 includes a weight value indicating the relative likelihood that the particular transitions between modules will occur. In the example of Table 1, higher weight values indicate less likely transitions. The server 401 may apply a shortest-path graph traversal algorithm (also known as a "least cost" algorithm) to determine a desirable module streaming sequence based on the currently executing module. Example shortest-path algorithms may be found in *Telecommunications Networks: Protocols, Modeling and Analysis*, Mischa Schwartz, Addison Wesley, 1987, § 6.

Table 1: Graph Edge Table

Edge	Weight
(A,B)	1
(A,D)	7
(A,E)	3
(B,C)	1
(B,G)	3
(C,E)	2
(C,G)	6

(D,F)	2
(E,H)	2
(F,H)	1
(G,E)	3
(G,H)	4

For example, Table 2 represents the minimum path weight between module "A" and the remaining modules of Table 1.

**Table 2: Shortest Paths from Application Module "A"**

<u>From</u>	<u>To</u>	<u>Shortest Path Weight</u>	<u>Path</u>
A	B	1	A-B
	C	2	A-B-C
	D	7	A-D
	E	3	A-E
	F	9	A-D-F
	G	4	A-B-G
	H	5	A-E-H

Based on the weight values shown in Table 2, the server 401 may determine that, during the execution of module "A", the module streaming sequence "B," "C," "E," "G," "H," "D," "F" is advantageous. If a particular module in a determined sequence is already present at the client 402, as may have been reported by control data 415, the server 401 may eliminate that module from the stream of modules 405. If, during the transmission of the sequence "B," "C," "E," "G," "H," "D," "F," execution of module "A" completes and execution of another module begins, the server may interrupt the delivery of the sequence "B," "C," "E," "G," "H," "D," "F," calculate a new sequence based on the now executing module, and resume streaming based on the newly calculated streaming sequence. For example, if execution transitions to module "B" from module "A," control data 415 may be sent from the client 410 to the server 401 indicating that module "B" is the currently executing module. If module "B" is not already available at the client 410, the server 401 will complete delivery of module "B" to the client and determine a new module streaming sequence. By applying a shortest-path routing algorithm to the edges of Table 1 based on module "B" as the starting point, the

minimum path weights between module "B" and other modules of the graph 300 (Fig. 3) can be determined, as shown in Table 3.

**Table 3: Shortest Paths from Module "B"**

<u>From</u>	<u>To</u>	<u>Shortest Path Weight</u>	<u>Path</u>
B	C	1	B-C
	E	5	B-C-E
	G	3	B-G
	H	7	B-C-E-H

Based on the shortest path weights shown in Table 3, the server 401 may determine that  
5 module streaming sequence "C," "G," "E," and "H" is advantageous.

Other algorithms may also be used to determine a module streaming sequence. For example, a weighted graph 300 may be used wherein heavier weighted edges indicate a preferred path among modules represented in the graph. In Table 4, higher assigned weight values indicate preferred transitions between modules. For example, edges (A,B), (A,D), and  
10 (A,E) are three possible transitions from module A. Since edge (A,B) has a higher weight value than edges (A,D) and (A,E) it is favored and therefore, given module "A" as a starting point, streaming of module "B" before modules "D" or "E" may be preferred. Edge weight values can be, for example, a historical count of the number of times that a particular module was requested by a client, the relative transmission time of the code module, or a value  
15 empirically determined by a system administrator and stored in a table 402 at the server 401. Other edge weight calculation methods may also be used.

**Table 4: Preferred Path Table**

<u>Edge</u>	<u>Weight</u>
(A,B)	100
(A,D)	15
(A,E)	35
(B,C)	100
(B,G)	35
(C,E)	50

(C,G)	20
(D,F)	50
(E,H)	50
(F,H)	100
(G,E)	35
(G,H)	25

In an preferred-path (heavy weighted edge first) implementation, edges in the graph 300 having higher weight values are favored. The following exemplary algorithm may be used to determine a module streaming sequence in a preferred-path implementation:

- 1: Create two empty ordered sets:
  - i) A candidate set storing pairs (S,W) wherein "S" is a node identifier and "W" is a weight of an edge that may be traversed to reach node "S."
  - ii) A stream set to store a determined stream of code modules.
- 2: Let  $S_i$  be the starting node.
- 3: Append the node  $S_i$  to the Stream Set and remove any pair  $(S_i, W)$  from the candidate set.
- 4: For each node  $S_j$  that may be reached from node  $S_i$  by an edge  $(S_i, S_j)$  having weight  $W_j$ :
  - If  $S_j$  is not a member of the stream set then add the pair  $(S_j, W_j)$  to the candidate set.
  - If  $S_j$  appears in more than one pair in the candidate set, remove all but the greatest-weight  $(S_j, W)$  pair from the candidate set.
- 5: If the Candidate set is not empty
  - Select the greatest weight pair  $(S_k, W_k)$  from the candidate set.
  - Let  $S_i = S_k$
  - Repeat at step 3

For example, as shown in Table 5, starting at node "A" and applying the foregoing algorithm to the edges of Table 4 produces the stream set {A, B, C, E, H, G, D, F}:

Table 5: Calculation of Stream Set

Iteration	{Stream Set} / {Candidate Set}
1	{A} / {(B,100) (D,15) (E,35)}
2	{A, B} / {(D,15) (E,35) (C,100) (G,35)}
3	{A, B, C} / {(D,15) (E,35) (G,35)}
4	{A, B, C, E} / {(D,15) (G,35) (H,50)}
5	{A, B, C, E, H} / {(D,15) (G,35)}
6	{A, B, C, E, H, G} / {(D,15)}
7	{A, B, C, E, H, G, D} / {(F,50)}
8	{A, B, C, E, H, G, D, F} / {}

Implementations may select alternative algorithms to calculate stream sets.

Application streaming may also be used to stream subsections of an application or module. For example, subsections of compiled applications, such as applications written in 5 C, C++, Fortran, Pascal, or Assembly language may be streamed from a server 401 to a client 410. Referring to Fig. 5A, an application 500 may include multiple code modules such as a main code module 501 and code libraries 510 and 515. The main module 501 contains program code that is executed when the application is started. The code libraries 510 and 515 may contain header data 511 and 516 as well as executable procedures 512-514 and 517-519 10 that are directly or indirectly called from the main module 501 and other library procedures.

In a Microsoft Windows 95 / Microsoft Visual C++ implementation, the main code module 501 may contain a compiled C++ "main" procedure and the library modules 510 and 515 may be dynamic link libraries having compiled C++ object code procedures. Header data 511 and 516 may include symbolic names used by operating system link procedures to 15 dynamically link libraries 510 and 515 with the main module 501. Header data may also indicate the location of each procedure within the library. In a jump table implementation, a calling procedure may access library procedures 512-514, 517-519 by jumping to a predetermined location in the header 511 or 516 and from there, accessing additional code and/or data resulting in a subsequent jump to the start of the procedure.

20 Data and procedures within an application's code modules and libraries may be many hundreds or thousands of bytes long. Prior to executing an application, a client may need to retrieve a lengthy set of modules and libraries. By reducing the size of the module and library set, the initial delay experienced prior to application execution can be reduced. In a streaming

implementation of application 500, code within subsections of the application's code modules can be removed and replaced by shortened streaming "stub" procedures. The replacement of application code with streaming stub procedures may reduce module size and associated transmission delay. For example, referring to Figs. 5A and 5B, the code library 510 may include a header 511 that is 4 kilobytes (Kbytes) in length and procedures 512-514 that are, respectively, 32 Kbytes, 16 Kbytes, and 8 Kbytes. Referring to Figs. 5B and 5C, to reduce the size of the library 510, procedures code 512-514 may be removed from the library 510 and stored in a streaming code module database 403 at the server 401 (Fig. 4). The removed procedure code 512-514 may be replaced by "stub" procedures 515-517 resulting in a reduced-size code library 530 that can be linked with application modules 501 and 520 in place of library 510. Header data 511 of library 530 can include updated jump or link information allowing stub procedures 515-517 to act as link-time substitutes for procedures 512-514.

A server 401 may provide a streaming-enabled version of application 500 to a client 410 by sending main module 501, library module 520, "streamed" library 530, and, in some implementations, a streaming support file 535 to the client 410 in response to a request for the application 500. The streaming support file 535 may include procedures accessed by the stubs 515-517 to facilitate code streaming between the server 401 and client 410. At the client 410, modules 501, 520, 530 and 535 can be linked and execution of the resulting application can begin. As the main module 501 and various called procedures are executed at the client 410, code modules stored in the database 403 can be streamed from the server 401 to the client 410. Data may be included in the stream 403 to identify stub procedures 414-417 associated with the streamed code modules. As the streamed modules are received at the client, they are integrated with the executing application.

In an appended module implementation, streamed code modules are integrated with the executing application by appending received modules to their corresponding library or code file. For example, referring to Figs. 5C and 5D, as modules 512-514 are streamed from the server to the client, they are appended to the library file 530 thereby forming an augmented library file 540. As the modules 512-514 are streamed from the server 401 and appended to the file 530, header data 511 or stub data 515-516 is updated so that the now-appended modules are accessible from a calling procedure. For example, referring to Fig. 5D, an additional "jump" may be added between each stub procedure 515-517 and its associated

appended module 512-514. Alternatively, header data 511 may be updated so that procedures 512-514 are accessible in place of stubs 515-517. In a stub-replacement implementation, stubs 515-516 are replaced by procedure modules 512-514 as the modules are received from the server 401. Stub replacement may require altering or rearranging the location of the 5 remaining stubs or procedures within a code module or library as replacement code is received. Implementations may employ still other methods of integrating streamed code with executing applications and modules.

In some scenarios, removed code, such as procedure code 512-514 which, in the example given, was replaced by stubs 515-517, may be required (called by another 10 procedure) before it is streamed from the server 401 and integrated with the module 530. In such a case, stub code 515-516 may access streaming functions in the streaming support library 535 to obtain the required procedure. To do so, the streaming support library 535 may send control data 415 to the server 401 to request the needed procedure. In response, the server 401 can halt the current module stream 405 and send the requested module. Upon 15 receipt of the requested module, procedures in the streaming support library 535 may be used to integrate the received module with the application and to continue with the execution of the requested module. The server may thereafter determine a new module stream based on the requested module or other control data 415 that was received from the client.

Code modules may be reduced in size without the use of stub procedures. For 20 example, referring again to Figs. 4, 5A, 5B, and 5E, in a interrupt driven implementation, procedure code 512-514 may be removed from a code library 510 and stored in a database 403. Header information 511 as well as data indicating the size and location of removed procedure code 512-514 may then be transmitted to a client 410. The client 410 may construct a new library 550 by appending a series of interrupt statements in place of the 25 removed procedure code 512-514. When the application 500 is executed, the code library 550 is substituted for the library 510 and execution of the program 500 may begin. As the program 500 executes, the removed procedure code 512-514 can be streamed to the client 410 and stored in a local database 411. If the application 500 attempts to execute procedure code 512-514 it may instead execute one of the interrupt statement that have replaced 30 procedure code 512-514. The execution of the interrupt statement halts the execution of the program 500 and transfers control to a streaming executor program 415.

Executor 415 implements interface technology similar to that of a conventional runtime object code debugger thereby allowing the executor 415 to intercept and process the interrupt generated by the application 500. When the interrupt is intercepted by the executor 415, data provided to the executor 415 as part of the client execution platform (operating system) interrupt handling functionality can be used to identify the module 550 in which the interrupt was executed and the address of the interrupt code within the module. The executor 415 then determines whether procedure code 512-514 associated with the interrupt location has been received as part of the module stream 415 sent to the client. If the appropriate procedure code has been received, the executor 515 replaces the identified interrupt with the 5 its corresponding code. For example, procedures 512-514 may be segmented into 4 Kilobyte code modules that are streamed to the client 410. When an interrupt statement is executed by the application 500, the executor 415 intercepts the interrupt, determines an appropriate 4 Kilobyte code block that includes the interrupt statement, and replaces the determined code 10 block with a received code module. If the appropriate code module has not yet been received, an explicit request may be sent from the client 410 to the server 401 to retrieve the code 15 module prior to its insertion in the library 550. The executor 415 may thereafter cause the application 500 to resume at the address of the encountered interrupt.

Implementations may also stream entire modules or libraries. For example, main code module 501 may be received from the server 401 and begin execution at the client 410 while 20 code libraries 510 and 515 are streamed from the server 401 to the client 410. Integration of streamed modules with executing modules may be provided by client 410 dynamic module linking facilities. For example, delay import loading provided by Microsoft Visual C++ 6.0 may be used to integrate streamed modules 510 and 515 with executing modules 501.

Dynamic linking of streamed modules may be facilitated by storing the streamed modules on 25 a local hard disk drive or other storage location accessible by client 410 link loading facilities. In an exemplary implementation, streaming is facilitated by altering client 410 operating system link facilities such that the link facility can send control data 415 to the server 401 to request a particular module if the module is has not already been streamed to the client 401.

30 In a protected-memory computer system, direct manipulation of executing application code and data may be restricted. In such systems, a "kernel" level processes or procedure may be required to support integration of streamed modules with executing application. In

such a case, streaming support 535 may be pre-provisioned by installing support procedures at the client 410 prior to the client's request for the application 500.

Other methods of determining stream sets may be used. In a list-based implementation, the streaming control file may include predetermined list of module streaming sequences. For example, the streaming control file 402 may include a module streaming sequence list associated with a first user and a second module streaming sequence list associated with a second user. Control data 415 sent from the client 410 to the server 401 may identify the current user at the client 410. Once the user has been identified to the server, the server may stream software modules in accordance with the user's associated streaming sequence list. User-based streaming data may be advantageous where a user's past behavior can be used to anticipate the order of modules to be accessed by that user.

In graph-based streaming control file implementations, the weight of edges connecting nodes may be determined statically or dynamically and may be determined based on a collection of historical usage data. For example, in a programmer-controlled implementation, a software programmer estimate the likelihood that particular transitions between nodes will occur based on the programmer's knowledge of the software code and the expected application usage patterns. Alternatively, application profiling programs may be used to gather run-time execution data recording transitions between various applets, Classes or code modules and thereby determine the likelihood that particular transitions will occur. In a client-feedback implementation, control data 415 sent from the client 410 to the server 401 during module execution is used to build a statistical database of module usage and, based on that database, determine the module streaming order.

In a client-controlled streaming implementation, streaming control data 402 may be located at the client 410 and control data 415 sent from the client 410 to the server 401 may be used to sequentially request a stream of modules from the server. For example, while the client computer 410 is executing a first module, a background process may send control data 415 to a server to request additional modules that can be buffered on a hard disk 411 at the client computer 410. A client-controlled streaming implementation may used existing HTTP servers and HTTP protocols to send request from the client 410 to the server 401 and send software modules from the server 401 to the client 410. Furthermore, although streaming of software modules has been emphasized in the foregoing description, non-executable data,

such as hypertext markup language, binary graphic files, and text, may be streamed as a collection of modules.

5        Implementations may include a "handshaking" procedure whereby, at the start of application execution, control data 415 is sent between the server 401 and the client 410. The handshaking data may include an inventory of application modules residing at the client and at the server. Such handshaking data allows both the client 410 and server 401 to determine their respective software module inventory and to optimize the stream of software modules based on that inventory information.

10      In a history-dependent implementation, a server or client can store data about a series of transitions between modules and calculate a new module stream based on a history of transitions. For example, referring to Fig. 3, if the module "G" was reached by the path A-B-G, then a server or client may determine that module "E" followed by "H" is to be streamed. On the other hand, if the module "G" was reached by the path A-B-C-G then the streaming sequence may include only the module "H."

15      Locally stored data and data at remote locations may be accessed at a client computer using a data browser application. For example, Netscape Navigator® and Microsoft Internet Explorer® are browser applications can display visual "links" that are selectable by a user to access data. When a link is selected by a user, data associated with the link can be retrieved from a server and presented to the user. Different data types may be presented in different 20 ways. For example, a Java applet may be presented as an executing program, encoded audio data may be presented as sounds reproduced by the client computer, and text may be presented visually on a graphic display.

25      Data identified by a link to a remote computer system can be streamed from the remote system to a client computer prior to a user's request for the data. Data may also be obtained from remote computer systems by a caching program at the client computer that automatically retrieves data associated with links in an HTML page prior to a user's request for the data. Such streamed or autonomously cached data can be stored at a client computer in local storage space. When the link to the remote source of such data is selected, the locally stored copy of the data can be presented to a user, thereby improving the access speed to the 30 data.

Advantages can be obtained by distinctly presenting links to data that has been streamed, autonomously cached, or otherwise locally stored at a client computer prior to a

user's request for the data. For example, by distinctly presenting such links, a user may be made aware of data that has not been viewed but which may be rapidly accessed. Such an awareness may increase the likelihood that a user will access the link to obtain data that has not been previously presented. In server-controlled streaming implementations, a streaming control database at a server can give high streaming priority to data associated with preferred links, such as links to particular advertisers. Enabling the distinct presentation of such preferred links to streamed data may provide the advantage of increasing the likelihood that such links will be selected by a user.

Different distinguishing characteristics may be used to distinctly present links to data that has been streamed, cached, or otherwise locally stored at a client computer prior to a user's request for the data. For example, as shown in Fig. 6, a browser application 600 can use a normal typeface to display an ordinary link 610, a bold-faced and italicized typeface to display a link that has been previously accessed by a user 620, and an underlined typeface to display a link 630 to data locally stored at a client computer prior to a user's request for the data. A user of the browser 600 will thereby know that unseen data associated with the link 630 can be presented with a low delay.

A determination of link presentation characteristics can be made by a browser prior to displaying link information to a user. Referring to Fig. 7, in an exemplary browser application, a link history file may be examined to determine 702 whether data associated with a link has been recently viewed. If it is determined that data associated with the link has been recently viewed, a set of characteristics identifying the link as a recently viewed link 704 is associated with the link and used when displaying the link to a user 707.

Alternatively, if the link's associated data has not been recently viewed, but is stored locally at the client computer, a set of characteristics indicating that the data is rapidly accessible may be assigned to the link 705 and used when displaying the link 707. Finally, if other presentation characteristics are not assigned, a default set of characteristics 706 may be assigned to the link and used for presentation of the link 707. A link's presentation may be modified when certain events occur at the client computer. For example, the browser may detect activity 708 such as link selection 709 by a user or a local storage change 710. The local storage change 710 may occur, for example, when data is streamed from a remote computer or obtained from a remote computer by autonomous caching software prior to a user's request to access the data. After a link selection 709 (and the presentation of that link's

associated data 711) or after a storage change 711 activity is detected, the browser may repeat link characteristics determination and link presentation 701.

The invention may be implemented in computer hardware, firmware, software, digital electronic circuitry or in combinations of them. Apparatus of the invention may be

- 5    implemented in a computer program product tangibly embodied in a machine-readable storage device for execution by a programmable processor; and method steps of the invention may be performed by a programmable processor executing a program of instructions to perform functions of the invention by operating on input data and generating output.

The invention may advantageously be implemented in one or more computer  
10    programs that are executable on a programmable system including at least one programmable processor coupled to receive data and instructions from, and to transmit data and instructions to, a data storage system, at least one input device, and at least one output device. Each computer program may be implemented in a high-level procedural or object-oriented programming language, or in assembly or machine language if desired; and in any case, the  
15    language may be a compiled or interpreted language. Suitable processors include, by way of example, both general and special purpose microprocessors. Generally, a processor will receive instructions and data from a read-only memory and/or a random access memory.

Storage devices suitable for tangibly embodying computer program instructions and data include all forms of non-volatile memory, including by way of example semiconductor  
20    memory devices, such as EPROM, EEPROM, and flash memory devices; magnetic disks such as internal hard disks and removable disks; magneto-optical disks; and CD-ROM disks. Any of the foregoing may be supplemented by, or incorporated in, specially-designed ASICs (application-specific integrated circuits).

A number of embodiments of the present invention have been described.  
25    Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. For example, links can be distinguished by characteristics other than, or in addition to, typeface. For example, in a graphical user interface (GUI) environment, a mouse-controlled pointer may assume a distinct shape when it is in the proximity of a link to streamed data or a particular sound may be heard when the  
30    pointer is in the proximity of the link, a unique color may be used, or a particular shape may be associated with the link. Accordingly, other embodiments are within the scope of the following claims.

**WHAT IS CLAIMED IS:**

1. A computer-implemented data access method comprising:
  - 2 receiving data at a computer from another computer prior to a user request for the data;
  - 3 storing the received data at the computer; and
  - 4 presenting, prior to a user request for the data,[JVM1] a link associated with the received data using a distinguishing characteristic that differentiates the link from links associated with data not stored at the computer[JVM2].
2. The method of claim 1 wherein:
  - 3 the other computer comprises data storage comprising a stored collection of executable modules associated with an application;
  - 4 the data received at the computer from the other computer comprises a module selected from the stored collection; and
  - 5 the computer provides an execution environment for executing modules of the application.
3. The method of claim 2 further comprising:
  - 4 forming a module set comprising the data by selecting a sequence of modules from the collection in accordance with predetermined criteria; and
  - 5 streaming the module set to the computer from the other computer.
4. The method of claim 3 wherein selecting a sequence comprises selecting in accordance with a streaming control database at the other computer.
5. The method of claim 1 further comprising:
  - 2 receiving other data at the computer in response to a user request for the other data;
  - 3 storing the received other data at the computer; and
  - 4 presenting a link associated with the other data using a different distinguishing characteristic.

- 1    6. The method of claim 5 further comprising presenting a link to data not stored at the
- 2       computer using another different distinguishing characteristic.
  
- 1    7. The method of claim 6 wherein the distinguishing characteristic comprises a font and the
- 2       different distinguish characteristic comprises a different font.
  
- 1    8. The method of claim 6 wherein the distinguishing characteristic comprises a graphical
- 2       user interface pointer shape and the different distinguish characteristic comprises a
- 3       different graphical user interface pointer shape.
  
- 1    9. The method of claim 1 wherein the link comprises a pointer to the data at the other
- 2       computer.
  
- 1    10. The method of claim 9 further comprising:
  - 2       receiving a user request comprising a selection of the link; and
  - 3       presenting the stored received data in response to the user request.
  
- 1    11. The method of claim 10 wherein the received data comprises a program module and
- 2       presenting the received data comprises executing the program module.
  
- 1    12. The method of claim 10 where the received data comprises text and presenting the stored
- 2       received data comprises displaying the text on a video display.
  
- 1    13. The method of claim 10 further comprising altering the presentation of the link after
- 2       presenting the stored received data.
  
- 1    14. A computer program residing on a computer-readable medium, comprising instructions
- 2       for causing a computer to:
  - 3       receive data from another computer prior to a user request for the data;
  - 4       store the received data; and

5 present, prior to a user request for the data, a link associated with the received data using  
6 a distinguishing characteristic that differentiates the link from links associated with  
7 data not stored at the computer.

1 15. The program of claim 14 wherein:

2 the other computer comprises data storage comprising a stored collection of executable  
3 modules associated with an application;  
4 the data received at the computer from the other computer comprises a module selected  
5 from the stored collection; and  
6 the program residing on the computer-readable medium further comprises instructions to  
7 provide an execution environment for executing modules of the application.

1 16. The program of claim 15 further comprising instructions for causing the computer to:

2 select a sequence of modules from the collection in accordance with predetermined  
3 criteria; and  
4 send data instructing the other computer to send the sequence of modules.

1 17. The method of claim 16 wherein the instructions to select a sequence of modules  
2 comprise instructions to access the predetermined criteria from a streaming control  
3 database to compute the sequence.

1 18. The program of claim 14 further comprising instructions for causing the computer to:

2 receive other data in response to a user request for the other data;  
3 store the received other data at the computer; and  
4 present a link associated with the other data using a different distinguishing  
5 characteristic.

1 19. The program of claim 18 further comprising instructions for causing the computer to  
2 present a link to data not stored at the computer using another different distinguishing  
3 characteristic.

- 1    20. The program of claim 14 wherein the link comprises a pointer to the data at the other
- 2       computer.
- 1    21. The program of claim 20 further comprising instructions for causing the computer to:
  - 2       receive a user request comprising a selection of the link; and
  - 3       present the stored received data in response to the user request.
- 1    22. The program of claim 21 further comprising instructions for causing the computer to alter  
2       the presentation of the link after presenting the stored received data.
- 1    23. A computer system comprising:
  - 2       an interface operatively coupled to a communications link to exchange data with another
  - 3       computer;
  - 4       a display device for the presentation of data to a user;
  - 5       an input device to receive a user request;
  - 6       a memory storing executable instructions; and
  - 7       a processor operatively coupled to the interface, the image display, the input device, and
  - 8       the memory, the processor being configured by the stored executable instructions to:
    - 9       (a) receive data over the interface from another computer prior to a receipt of a
    - 10      request at the input device for the data;
    - 11      (b) store the received data; and
    - 12      (c) present a link on the display device prior to a user request for the data, the link
    - 13      being associated with the received data and presented using a distinguishing
    - 14      characteristic that differentiates the link from links associated with data not stored
    - 15      at the computer.
- 1    24. The system of claim 23 wherein the stored executable instructions further comprise  
2       instructions to configure the processor to:
  - 3       receive other data in response to a user request for the other data received at the input
  - 4       device;

5 store the received other data at the computer; and  
6 present another link on the display device, the other link being associated with the other  
7 data and presented using a different distinguishing characteristic.

1 25. In a computer system having a processor operatively interconnected to a memory, a  
2 network communications device, a graphical display device, a data storage device, and a  
3 user input device, a graphical user interface comprising:  
4 a plurality of interface objects each user-selectable to present data associated with the  
5 object to a user; and  
6 a plurality of visual characteristics associated with interface objects, the visual  
7 characteristics comprising a characteristic associated with objects having associated  
8 data stored locally at the computer system and not previously presented to the user  
9 and a different visual characteristic associated with objects having associated data  
10 stored at the computer system which has been previously presented to the user.

1 26. The graphical user interface of claim 25 wherein data stored locally at the computer  
2 system comprises data received from another computer system preceding a user's  
3 selection of an interface object associated with the data from the other computer system.

1 27. The graphical user interface of claim 25 wherein each interface object is a hypertext  
2 markup language link.




Fig. 1  
(Prior Art)

BEST AVAILABLE COPY




Fig. 2




Fig. 3




Fig. 4




Fig. 5D




Fig. 5C




Fig. 5B




Fig. 5A




Fig. 5E




Fig. 6

BEST AVAILABLE COPY

9/9




Fig. 7

# INTERNATIONAL SEARCH REPORT

International Application No  
PCT/US 00/02190

<b>A. CLASSIFICATION OF SUBJECT MATTER</b> IPC 7 G06F17/30		
According to International Patent Classification (IPC) or to both national classification and IPC		
<b>B. FIELDS SEARCHED</b> Minimum documentation searched (classification system followed by classification symbols) IPC 7 G06F		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practical, search terms used)		
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	ZHIMEI JIANG ET AL: "Prefetching links on the WWW" IEEE INTERNATIONAL CONFERENCE ON COMMUNICATIONS (ICC), US, NEW YORK, NY: IEEE, 1997; pages 483-489, XP002086568 ISBN: 0-7803-3926-6 page 488, column 2, paragraph 4 - paragraph 5; figures 6,7 page 483, column 1, paragraph 2	1-4, 9-17, 20-23, 25-27
X	WO 97 46955 A (AT & T CORP) 11 December 1997 (1997-12-11) page 3, line 16 - line 24 abstract	1-6, 9, 10, 12-27
A	US 5 742 768 A (ZAMORA BENJAMIN A ET AL) 21 April 1998 (1998-04-21) abstract	2
-/-		
<input checked="" type="checkbox"/> Further documents are listed in continuation of box C.		<input checked="" type="checkbox"/> Patent family members are listed in annex.
<b>* Special categories of cited documents :</b>		
*A* document defining the general state of the art which is not considered to be of particular relevance *E* earlier document but published on or after the international filing date *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) *O* document referring to an oral disclosure, use, exhibition or other means *P* document published prior to the international filing date but later than the priority date claimed		
*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art *S* document member of the same patent family		
Date of the actual completion of the international search  20 June 2000		Date of mailing of the international search report  05/07/2000
Name and mailing address of the ISA European Patent Office, P.B. 5818 Patentdaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl. Fax (+31-70) 340-3016		Authorized officer  Triest, J

## INTERNATIONAL SEARCH REPORT

International Application No  
PCT/US 00/02190

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>CHEN, Z. ET AL: "Real Time Video and Audio in the World Wide Web"          FOURTH INTERNATIONAL WORLD WIDE WEB CONFERENCE, 'Online!          11 - 14 December 1995, pages 1-14,          XP002107740          BOSTON, MASSACHUSETTS, USA          Retrieved from the Internet:          &lt;URL:<a href="http://www.w3.org/Conferences/WWW4/Papers/211">http://www.w3.org/Conferences/WWW4/Papers/211</a>&gt; 'retrieved on 1999-06-29!          paragraph '06.2!</p> <hr/>	3,4
P,X	<p>US 5 878 223 A (MCLEAN JAMES GORDON ET AL) 2 March 1999 (1999-03-02)          abstract</p> <hr/>	1,9-17, 20-23, 25-27

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No  
PCT/US 00/02190

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
WO 9746955 A	11-12-1997	US 5870769 A		09-02-1999
		CA 2256722 A		11-12-1997
		EP 0902926 A		24-03-1999
US 5742768 A	21-04-1998	NONE		
US 5878223 A	02-03-1999	NONE		